

PD-Quant: Post-Training Quantization Based on Prediction Difference Metric

Jiawei Liu^{1,*, \diamond} Lin Niu^{1,*, \diamond} Zhihang Yuan^{2, \dagger} Dawei Yang² Xinggang Wang¹ Wenyu Liu^{1, \dagger}
¹ School of EIC, Huazhong University of Science & Technology ² Houmo AI
 {jiaweiliu, linniu}@hust.edu.cn hahnyuan@gmail.com dawei.yang@houmo.ai
 {xgwang, liuwu}@hust.edu.cn

Abstract

Post-training quantization (PTQ) is a neural network compression technique that converts a full-precision model into a quantized model using lower-precision data types. Although it can help reduce the size and computational cost of deep neural networks, it can also introduce quantization noise and reduce prediction accuracy, especially in extremely low-bit settings. How to determine the appropriate quantization parameters (e.g., scaling factors and rounding of weights) is the main problem facing now. Existing methods attempt to determine these parameters by minimize the distance between features before and after quantization, but such an approach only considers local information and may not result in the most optimal quantization parameters. We analyze this issue and propose PD-Quant, a method that addresses this limitation by considering global information. It determines the quantization parameters by using the information of differences between network prediction before and after quantization. In addition, PD-Quant can alleviate the overfitting problem in PTQ caused by the small number of calibration sets by adjusting the distribution of activations. Experiments show that PD-Quant leads to better quantization parameters and improves the prediction accuracy of quantized models, especially in low-bit settings. For example, PD-Quant pushes the accuracy of ResNet-18 up to 53.14% and RegNetX-600MF up to 40.67% in weight 2-bit activation 2-bit. The code is released at <https://github.com/hustvl/PD-Quant>.

1. Introduction

Various neural networks have been used in many real-world applications with high prediction accuracy. When deployed on resource-limited devices, networks' vast memory and computation costs become significant challenges. Re-

ducing overhead while maintaining the model accuracy has received considerable attention. Network quantization is an effective technique that can compress the neural networks by converting the format of values from floating-point to low-bit [10, 12, 27]. There are two types of quantization: post-training quantization (PTQ) [32] and quantization-aware training (QAT) [18]. QAT requires retraining a model on the labeled training dataset, which is time-consuming and computationally expensive. While PTQ only requires a small number of unlabeled calibration samples to quantize the pre-trained models without retraining, which is suitable for quick deployment. Existing PTQ methods can achieve good prediction accuracy with 8-bit or 4-bit quantization by selecting appropriate quantization parameters. [22, 23, 30]. Local metrics (such as MSE [7] or cosine distance [45] of the activation before and after quantization in layers) are commonly used to search for quantization scaling factors. These factors are chosen layer by layer by minimizing the local metric with a small number of calibration samples. In this paper, we observe that there is a gap between the selected scaling factors and the optimal scaling factors .

Since the noise from quantization will be more severe at low-bit, the prediction accuracy of the quantized model significantly decreases at 2-bit. Recently, some methods [24, 25, 44] have added a new class of quantization parameters, weight rounding value, to adjust the rounding of weights. They optimize both quantization scaling factors and rounding values by reconstructing features layer-wisely or block-wisely. Besides, the quantized model by PTQ reconstruction is more likely to be overfitting to the calibration samples because adjusting the rounding of weights will significantly increase the PTQ's degree of freedom.

We propose an effective PTQ method, PD-Quant, to address the above-mentioned issues. In this paper, we focus on improving the performance of PTQ on extremely low bit-width. PD-Quant uses the metric that considers the

* Equal contribution. \diamond This work was done when Jiawei Liu and Lin Niu were interns at Houmo AI. \dagger Corresponding authors.

We define the optimal quantization scaling factors as the factors that make the quantized model have the lowest task loss (cross-entropy loss calculated by real label) on the validation set.

global information from the prediction difference between the quantized model and the full-precision (FP) model. We show that the quantization parameters optimized by prediction difference are more accurate in modeling the quantization noise. Besides, PD-Quant adjusts the activations for calibration in PTQ to mitigate the overfitting problem. The distribution of the activations is adjusted to meet the mean and variance saved in batch normalization layers. Experiments show that PD-Quant leads to better quantization parameters and improves the prediction accuracy of quantized models, especially in low-bit settings. Our PD-Quant achieves state-of-the-art performance in PTQ. For example, PD-Quant pushes the accuracy of weight 2-bit activation 2-bit ResNet-18 up to 53.14% and RegNetX-600MF up to 40.67%. Our contributions are summarized as follows:

1. We analyze the influence of different metrics and indicate that the widely used local metric can be improved further.
2. We propose to use the information of the prediction difference in PTQ, which improves the performance of the quantized model.
3. We propose Distribution Correction (DC) to adjust the activation distribution to approximate the mean and variance stored in the batch normalization layer, which mitigates the overfitting problem.

2. Related Work and Background

Many excellent works have been proposed to resolve neural networks’ enormous memory footprint and inference latency, including knowledge distillation [11, 17, 42], model pruning [28, 54], and model quantization [4, 35, 47].

We focus on model quantization, which is an effective technique for compressing neural networks. Quantized models keep their weights and activations in low-bit data types to reduce memory and computation requirements. We can map a real-valued tensor x (weights or activations) to the integer grid according to the following equation:

$$\tilde{x} = \text{clamp}\left(\left\lfloor \frac{x}{S} \right\rfloor + Z; q_{min}, q_{max}\right), \quad (1)$$

$$S = (x_{max} - x_{min}) / (2^b - 1), \quad (2)$$

where $\lfloor \cdot \rfloor$ is the round-to-nearest operator. S denotes the quantization scaling factors, which reflect the proportional relationship between FP values and integers. Moreover, Z is the offset defined as zero-point. x_{max} is the maximum in the vector, and x_{min} is the minimum in the vector. $[q_{min}, q_{max}]$ is the quantization range determined by

The prediction is the processed output of the last layer, such as the probability after softmax in the classification model.

bit-width. We only consider uniform unsigned symmetric quantization, as it is the most widely used quantization setup. Therefore, q_{min} is equal to 0 and q_{max} is equal to $2^b - 1$, where b is the bit-width which determines the number of integer grids. Nonuniform quantization [20] is challenging to deploy on hardware, so we will not consider it in this work. In general, we divide quantization methods into Quantization-aware training [5, 13, 21, 45, 49, 49, 53] (QAT) and Post-training quantization [3, 24, 29, 43, 44, 51] (PTQ).

2.1. Quantization-aware training

QAT modifies the quantization noise during training processes with full labeled training datasets. STE [2] solves the problem of producing zero gradients during backpropagation by employing a gradient estimator. During the training process, [19] smoothed activation ranges by exponential moving averages. LSQ [8] introduces trainable clipping threshold parameters to learn the min and max ranges of the quantization by STE. Although QAT enables lower bit quantization with competitive results, it needs labeled datasets and amounts of computing resources.

2.2. Post-training quantization

PTQ algorithms often determine the quantization scaling factors with limited calibration data through a simple parameter space search without training or fine-tuning [1, 7, 9, 41, 45, 51]. Metrics for searching scaling factors include MSE distance [7] and cosine distance [45]. [1] seeks the optimal clipping ranges by minimizing the difference between FP and quantized feature map.

Later, several methods have been proposed to optimize rounding values. Although these PTQ methods introduce fine-tuning during quantization, they still differ from QAT. QAT adjusts the model’s weights using whole labeled training dataset, while PTQ only optimizes the quantization parameters on some unlabeled data. AdaRound [29] suggests a new rounding mechanism that assigns a continuous variable to each weight value to determine whether it should be rounded up or down, rather than using the nearest rounding. BRECQ [24] optimizes activation scaling factors and proposes reconstructing quantization parameters block by block. QDrop [44] imports the activation error into the reconstruction process and introduces drop operation. These methods can achieve usable accuracy at low bits without inference overhead. They reconstruct features by calculating the MSE distance between quantized and FP activations. AQuant [26] improves activation quantization strategy and the quantization performance, but has additional inference overhead. NWQ [40] was published at the same time as our method, so we do not compare with it.

Some methods [3, 6, 14, 46, 50, 52] can now quantize models without using real data. GDFQ [46] adopts a generator to synthesize calibration data. ZeroQ [3] uses statis-

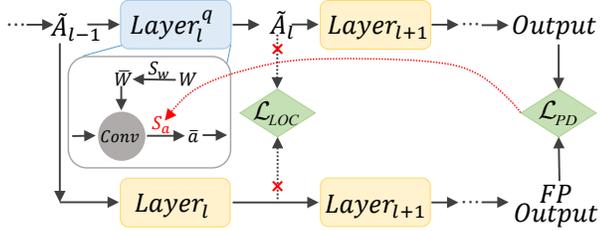


Figure 1. The pipeline overview of our method is with PD loss for determining activation scaling factors. The blue and yellow rectangle indicates the quantizing and FP layer, respectively. We mark the green diamond as the loss function. The red X in the figure represents an approach we did not adopt but adopted in previous work.

tics from batch normalization layers to synthesize calibration data from a Gaussian distribution. Compared to using real calibration data, these methods often struggle to exhibit competitive performance. However, inspired by ZeroQ, we proposed DC in PD-Quant.

3. Methodology

In this section, we will start by conducting a comprehensive analysis of the impacts of various metrics in the search for activation quantization scaling factors. Then, we develop a high-accuracy post-training quantization method, PD-Quant. It applies Prediction Difference (PD) loss to optimize the quantization parameters. Besides, we introduce regularization and propose Distribution Correction (DC) to solve the overfitting problem.

3.1. Prediction Difference Loss

Previous PTQ works [7, 45] search quantization scaling factors by local metrics, such as MSE or Cosine distance of each layer’s activation before and after quantization. To investigate the influence of these metrics when determining activation scaling factors, we compare their search results with those of task loss. Task loss refers to the cross-entropy loss determined by the real label, and we define the scaling factors with the lowest task loss as optimal. Fig. 2 shows the task loss and other metrics for different scaling factors of 2-bit activation. As seen, the scaling factor optimized by local metrics (Cosine and MSE) is inconsistent with that based on task loss. We can observe that the scaling factors searched by local metric losses is far from the optimal scaling factor minimized by task loss. As an example, as shown in Fig. 2a, for the activation quantization in the Resnet18.layer4.block0.conv1, green line of task loss (CE) indicates the optimal scaling factor is around $0.35 \times N_s$ (normalized scaling factor) while local metrics consider $0.15 \times N_s$ as their scaling factor.

Since the real label is not available in PTQ, we cannot

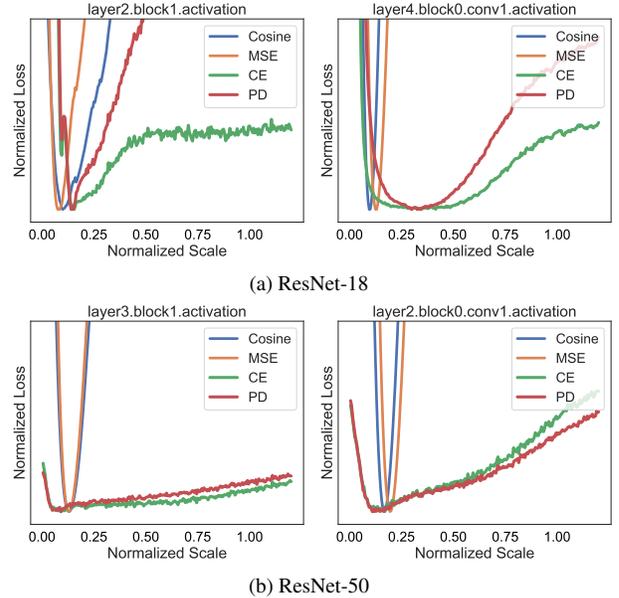


Figure 2. Different loss metrics are plotted, including task loss (CE) under various scaling factors on ResNet18 and ResNet50 with the ImageNet dataset. The normalized scaling factor (N_s) means the proportion of $\frac{x_{max} - x_{min}}{2^b - 1}$, and normalized loss indicates that the minimum value of all losses is normalized to 1.0.

use task loss as the metric. To achieve more accurate performance for PTQ, we propose PD loss to determine the activation scaling factors. PD loss compares the prediction difference between the FP and the quantized models rather than the difference in each layer’s activation before and after quantization. Specifically, we feed the quantized activation \tilde{A}_l into the subsequent FP layers to obtain the output prediction. And then, the output prediction will be compared with the FP output prediction to calculate PD loss. When searching for the optimal activation scaling factors, we take the following formula as the metric:

$$\arg \min_{S_a} \mathcal{L}_{PD}(f_l(\tilde{A}_{l-1}), f_{l+1}(L_l^q(\tilde{A}_{l-1}))), \quad (3)$$

where S_a denotes scaling factors for activation quantization

| Model | ResNet-18 | | ResNet-50 | | RegNet-600M | |
|----------------------|--------------|--------------|--------------|-------------|-------------|-------------|
| | W8A2 | W4A2 | W8A2 | W4A2 | W8A2 | W4A2 |
| Min-Max ^l | - | - | - | - | - | - |
| Cosine ^l | 11.09 | 4.15 | 2.19 | 1.14 | 0.96 | 0.65 |
| MSE ^l | 23.15 | 10.31 | 9.23 | 4.85 | 3.71 | 1.88 |
| PD ^g | 28.41 | 12.27 | 11.31 | 6.01 | 7.47 | 3.17 |

Table 1. Metric test (top-1 accuracy(%)) on validation set for activation scaling factors. The weight quantization for each bit setting in the table is the same. ^l represents calculating metric with only local information and ^g means metric with global information.

in quantizing layer L_l^q , and $f_l(\cdot)$ refers to a part of the FP network mapping input \tilde{A}_{l-1} to FP output. $f_{l+1}(\cdot)$ is the part of the FP network mapping the output of L_l^q to the final prediction. We make the assumption that the quantization for all previous layers of \tilde{A}_{l-1} has been done. Fig. 1 shows the detailed process for the parameters search by PD loss.

As shown in Fig. 2, the scaling factor optimized by the red line of PD loss (PD) is closer to the scaling factor expected by task loss. Still, take the ResNet-18 activation, as mentioned above, as an example. The optimal scaling factor chosen by PD loss is around $0.35 \times N_s$, almost the same as CE. In particular, we only consider the scaling factors of activation quantization, and we will discuss the scaling factors of weight quantization later in Sec. 5.1.

To further investigate how PD loss affects activation scaling factors determination, we respectively evaluate local metrics and PD loss on ResNet-18, ResNet-50 [15], and RegNetX-600MF [36] as examples. As shown in Tab. 1, the Min-Max quantization strategy (i.e. Eq. (1)) loses all quantization accuracy in the extremely low-bit. Compared to the local metric (Min-Max, MSE, Cosine), the quantization parameters optimization strategy with PD loss exhibits better results on ImageNet [37] dataset. The results verify that PD loss can significantly improve quantization performance in the extremely low-bit of the activation scaling factors. Moreover, it indicates that the widely used local metric could be better at low-bit. However, only optimizing the scaling factor does not gain usable results, although PD loss can achieve improvement. So we further verified the effect of our proposed PD loss with the optimization of both activation scaling factors and rounding values in Sec. 3.2, which can achieve usable accuracy at extremely low-bit.

We choose Kullback-Leibler (KL) divergence as the PD loss metric to measure the prediction difference for the classification networks. We will discuss the choice of PD loss in Sec. 5.2. For the detailed implementation of KL, we follow [17]. In summary, according to our analysis, considering prediction differences is beneficial for searching the activation scaling factors at low-bit.

3.2. Prediction Difference Loss with Regularization

Recently, some PTQ methods (such as AdaRound [29], BRECQ [24], and QDrop [44]) have achieved remarkable progress on PTQ, especially in low-bit. They propose optimizing a continuous variable for each weight value that will be quantized. These variables, called rounding values, determine whether weight values will be rounded up or down during the quantization process. The quantization parameters in these PTQ methods contain scaling factors and rounding values. In summary, we can describe the process

of weight quantization as follows:

$$\tilde{x} = clamp \left(\lfloor \frac{x + \theta}{S} \rfloor + Z; q_{min}, q_{max} \right), \quad (4)$$

where θ is the optimization variable for each weight value to decide rounding results up or down [29], i.e., $\frac{\theta}{S}$ ranges from 0 to 1. Other parameters are consistent with Eq. (1). As for activation quantization, the quantization steps are also the same as Eq. (1).

The next problem is whether PD loss (as described in Sec. 3.1) can also improve the optimization of both scaling factors and rounding values. We conduct preliminary experiments with PD loss on the ImageNet dataset with 1024 calibration samples to answer this question. We calculate PD loss as the difference between the prediction of the FP model and the prediction of the quantized model. However, as shown in Tab. 2, the performance of PD-only degrades severely in all settings. Moreover, it performs much worse on the validation set than on the calibration set. For example, the accuracy of 2 bit quantization (W2A2) on the calibration set is almost equal to the FP model on ResNet-18 [15], but the performance on the validation set is extremely poor. This phenomenon indicates that the model suffers from a severe overfitting problem. This issue arises from the limited amount of data available for calibration.

To address the above overfitting problem, we introduce an explicit regularization [16, 34, 44] by adding a constraint to the optimization problem. Our regularization approach constrains the difference between the FP and quantized activation of the internal block. On the one hand, the regularization can encourage the quantized model to learn the robustness from the FP model. On the other hand, regularization is beneficial to minimize the perturbation caused by quantization, which is verified to be effective in BRECQ [24]. Specifically, PD-Quant adopts PD loss to guide the quantization parameter optimization and introduces regularization to alleviate overfitting. In our implementation, we regard block as the smallest unit as in previous work [20, 24, 44, 48] and quantize neural networks from shallow to deep. In more

| Method | Bits (W/A) | Acc(val) | Acc(cali) |
|--------------------|------------|--------------|-----------|
| FP | 32/32 | 71.01 | 70.90 |
| PD-only | | 1.07 | 70.51 |
| PD+Reg | 2/2 | 49.16 | 71.09 |
| PD+Reg+Drop | | 52.74 | 68.26 |
| PD-only | | 51.32 | 70.41 |
| PD+Reg | 4/2 | 56.20 | 70.41 |
| PD+Reg+Drop | | 58.17 | 68.36 |

Table 2. PTQ accuracy on ImageNet at ResNet18 with 1024 calibration images. Reg means regularization. Acc(val)/Acc(cali) denotes the top-1 accuracy (%) on validation/calibration set.

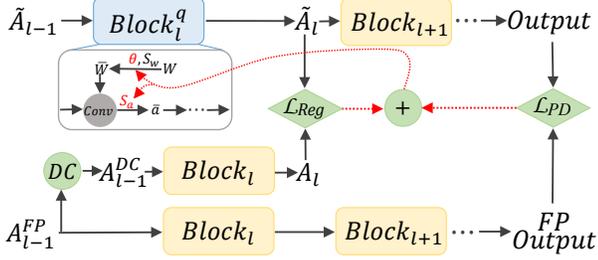


Figure 3. An overview of the PD-Quant. The blue and yellow rectangles indicate the quantized and FP layer, respectively. The green diamond is marked as the loss function. The green circle with DC indicates the Distribution Correction. FP output is the prediction of the whole FP network.

detail, when quantizing the l_{th} block, our optimization objective is as follows:

$$\arg \min_{\theta, S_a} \mathcal{L}_{PD}(O_{fp}, f_{l+1}(\tilde{A}_l) + \lambda_r \mathcal{L}_{reg}(A_l, \tilde{A}_l)), \quad (5)$$

$$\tilde{A}_l = B_l^q(\tilde{A}_{l-1}; \theta, S_a),$$

where θ is described in equation Eq. (4), and S_a is the activation quantization scaling factor. λ_r is a hyper-parameter to control the degree of regularization. \mathcal{L}_{reg} is the regularization loss to alleviate the overfitting problem. Here we use the MSE distance between each block’s output feature maps before and after quantization as $\mathcal{L}_{reg} = \|A_l - \tilde{A}_l\|_2^2$. B_l^q denotes the block being quantized with input \tilde{A}_{l-1} . Like Sec. 3.1, we also make the quantization for all previous layers of the input \tilde{A}_{l-1} has been done. $f_{l+1}(\cdot)$ is the part of the FP network mapping the output of quantizing block \tilde{A}_l to the final prediction. O_{fp} is the FP prediction as the target in PD loss. Fig. 3 shows the overview of our PD-Quant.

By introducing regularization, the performance of the quantized model has been dramatically improved, shown in Tab. 2 as PD+Reg. The overfitting problem has been effectively alleviated, and the gap in performance between the calibration set and the validation set has narrowed a lot. Random drop is also a regularization method, a supplement to our method for alleviating overfitting. We introduce activation drop in the feature map as QDROP [44], which can further improve the performance of the quantized model. The introduction of drop does not conflict with our method, and we only introduce it when computing the regularization loss.

3.3. Distribution Correction for Regularization

In this section, we will introduce a novel method to improve the generalizability of the quantized model further. Since only limited unlabeled images are accessible in PTQ, quantization parameters are determined only by activating these few samples. However, the feature distribution of such small data is difficult to reflect the feature distribution of the

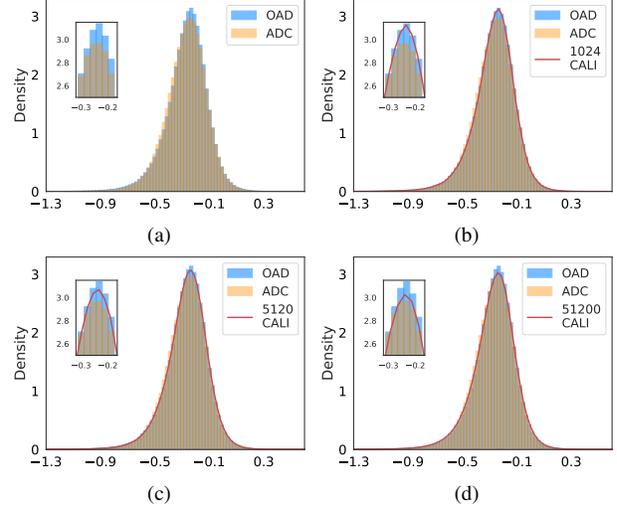


Figure 4. We plot the Distribution Correction for post-conv activation on the last block of ResNet18. OAD means the original activation distribution, and ADC denotes the original activation distribution (OAD) rectified by DC. Two histograms in blue and orange represent their distribution. The red line is the kernel density estimate (KDE) curve for different numbers of calibration (CALI) from the ImageNet training set.

whole training set. As described in Sec. 3.2, our regularization loss computes the distance between the quantized block activation and the FP block activation.

We propose a novel method to adjust the feature from the FP block itself, *i.e.*, rectifying the feature by the statistics stored in the pre-trained FP model’s Batch Normalization (BN) layer. Specifically, we correct the activation distribution of each FP block’s input by achieving an approximate mean and variance between the activation and BN statistics. Since these BN statistics come from the whole training set, the corrected distribution of activations will be closer to the distribution of the entire training set. Through learning the corrected FP feature distribution, the quantized model by DC has better generalizability.

Assuming there are n Batch Normalization (BN) layers in the l_{th} block, we can compute the mean and variance of their input, denoting $\{\hat{\mu}_{(i,l)}, \hat{\sigma}_{(i,l)} | i = 0, 1, \dots, n\}$. And the original mean $\mu_{(i,l)}$ and variance $\sigma_{(i,l)}$ are recorded in BN layers. We modify the input of the block A_{l-1}^{FP} to A_{l-1}^{DC} using the following optimization:

$$\arg \min_{A_{l-1}^{DC}} \lambda_c \sum_{i=1}^n (\|\hat{\mu}_{(i,l)} - \mu_{(i,l)}\|_2^2 + \|\hat{\sigma}_{(i,l)} - \sigma_{(i,l)}\|_2^2) + \|A_{l-1}^{DC} - A_{l-1}^{FP}\|_2^2, \quad (6)$$

where λ_c is a hyper-parameter. We fine-tuning the input A_{l-1}^{DC} so that $\hat{\mu}_{(i,l)}$ and $\hat{\sigma}_{(i,l)}$ are closer to $\mu_{(i,l)}$ and $\sigma_{(i,l)}$. The second term in Eq. (6) is to make the modified A_{l-1}^{DC}

| Methods | Bits (W/A) | ResNet-18 | ResNet-50 | MobileNetV2 | RegNetX-600MF | RegNetX-3.2GF | MNasx2 |
|-----------------|------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|
| Full Prec. | 32/32 | 71.01 | 76.63 | 72.62 | 73.52 | 78.46 | 76.52 |
| ACIQ-Mix [1] | | 67.00 | 73.80 | - | - | - | - |
| LAPQ [31] | | 60.30 | 70.00 | 49.70 | 57.71 | 55.89 | 65.32 |
| Bit-Split [43] | 4/4 | 67.56 | 73.71 | - | - | - | - |
| AdaRound [29] | | 67.96 | 73.88 | 61.52 | 68.20 | 73.85 | 68.86 |
| QDrop [44]* | | 69.17 | 75.15 | 68.07 | 70.91 | 76.40 | 72.81 |
| PD-Quant | | 69.23±0.06 | 75.16±0.07 | 68.19±0.12 | 70.95±0.12 | 76.65±0.09 | 73.26±0.09 |
| LAPQ | | 0.18 | 0.14 | 0.13 | 0.17 | 0.12 | 0.18 |
| Adaround | 2/4 | 0.11 | 0.12 | 0.15 | - | - | - |
| QDrop* | | 64.57 | 70.09 | 53.37 | 63.18 | 71.96 | 63.23 |
| PD-Quant | | 65.17±0.08 | 70.77±0.15 | 55.17±0.28 | 63.89±0.13 | 72.38±0.11 | 63.40±0.21 |
| QDrop* | 4/2 | 57.56 | 63.26 | 17.30 | 49.73 | 62.00 | 34.12 |
| PD-Quant | | 58.59±0.15 | 64.18±0.14 | 20.10±0.37 | 51.09±0.15 | 62.79±0.13 | 39.13±0.51 |
| QDrop* | 2/2 | 51.42 | 55.45 | 10.28 | 39.01 | 54.38 | 23.59 |
| PD-Quant | | 53.14±0.14 | 57.16±0.15 | 13.76±0.40 | 40.67±0.26 | 55.06±0.23 | 27.58±0.60 |

Table 3. Comparison on PD-Quant with various post-training quantization algorithms. * denotes our implementation using open-source codes. PD-Quant is our proposed method. Other results listed are all from [44]. We gain the results of 10 runs using randomly sampled calibration sets. The results in the table include the mean and standard deviation.

not deviate too much from A_{l-1}^{FP} . As shown in Fig. 3, we inference with A_{l-1}^{DC} to get A_l and calculate \mathcal{L}_{reg} .

In Fig. 4, we visualize the effect of our proposed DC. For all four subplots, the blue column (OAD) means an original activation distribution of 1024 calibration samples, and the orange column (ADC) denotes the OAD rectified by DC. Fig. 4a is the histogram of OAD and ADC. The red line in Fig. 4b is the kernel density estimate (KDE) curve for 1024 calibration samples. Next, we increase the number of calibration samples. Fig. 4c and Fig. 4d show the KDE curve for 5120 and 51200 calibration samples from the ImageNet training set. From Fig. 4, we notice that as the number of calibration samples increases, the KDE curve of calibration samples is closer to the ADC. The illustration shows that the distribution of ADC, produced by the DC from OAD, can reflect the distribution of more samples in the training set and improve the generalizability of the quantized model.

4. Experimental Results

4.1. Experimental Environments

We quantize various CNN architectures to evaluate our proposed method, including ResNet [15], MobileNetV2 [38], RegNet [36], and MNasNet [39]. The FP pre-trained models for all our implementation in the experiments are from [24]. Our method evaluates on ImageNet dataset [37] with a batch size of 32. We randomly sample 1024 images from the ImageNet training dataset as the calibration set. In addition, we also set the first and the last layer quantization to 8-bit for all PTQ experiments unless otherwise specified. We keep the same quantization settings and hyper-parameters in our implementation as QDrop [44].

The learning rate for the activation quantization scaling factor is $4e-5$, and for weight quantization rounding, the learning rate is $3e-3$. The DC is with a learning rate of $1e-3$. The choice of hyper-parameter λ_r and λ_c in Eq. (5) and Eq. (6) will be discussed later in Sec. 4.4, respectively. Quantization parameters are fine-tuning with 20000 iterations. Our code is based on Pytorch [33] and we evaluate all our experiments on a single Nvidia RTX A6000.

4.2. Performance Comparison

We comprehensively compare our PD-Quant with multiple PTQ algorithms in many bit-settings and find that PD-Quant achieves performance improvement, especially in extremely low-bit. Considering that only optimizing the activation scaling factors can not achieve usable results at low bits, all our experiments below optimize both rounding values and activation scaling factors unless otherwise specified. We choose the best-performing method QDrop [44] as our baseline. The results of QDrop in open-source code are higher than those in their paper. We introduced a drop [44] in the regularization process as described in Sec. 3.2. All our experiments with regularization include it unless otherwise stated.

In summarizing the results in Tab. 3, it can be observed that PD-Quant achieves significant improvements compared with those strong baselines of PTQ. When quantizing the network to W4A4, experiments show that PD-Quant slightly improves QDrop. However, the benefits of PD-Quant become more apparent as the bit-width decreases. At the W2A4 bit setting, the performance of PD-Quant is better than the baseline in all network architectures. With W4A2 quantization, PD-Quant can improve the

accuracy of MobileNetV2 by 2.8% and ResNet-600MF by 1.4%. For more challenging cases W2A2, the performance of PD-Quant surpasses QDrop in all networks. According to the table, PD-Quant reaches 27.58% in MNasNet, while QDrop only gets 23.59%. Since we use the same quantization setting with QDrop, the results indicate that the optimization strategy with our PD-Quant plays a critical role in extremely low-bit. Once PD-Quant finishes optimizing the quantization parameters, no additional computation is required for inference.

As in previous PTQ work, we also keep the first and last layer 8-bit. The results in Tab. 3 are all done with the 8-bit setting. Nevertheless, to be noted, some work [20, 24, 48] adds an extra first layer’s output 8-bit, which will perform better than our 8-bit setting. Therefore, we conduct further experiments to ensure the effectiveness of our method compared with these methods. Specifically, we use the same 8-bit settings as these methods for experiments, and the results are shown in Tab. 4.

| Methods | Bits (W/A) | ResNet-18 | MobileNetV2 |
|-----------------|------------|--------------|--------------|
| Full Prec. | 32/32 | 71.01 | 72.62 |
| BRECQ [24] | 4/4 | 69.60 | 66.57 |
| RAPQ [48] | | 69.28 | 64.48 |
| PD-Quant | | 69.72 | 68.76 |
| BRECQ [24] | 2/4 | 64.80 | 53.34 |
| RAPQ [48] | | 65.32 | 48.12 |
| PD-Quant | | 65.56 | 55.32 |

Table 4. Comparison on PD-Quant with different post-training quantization algorithms in another 8-bit setting.

4.3. Ablation Study

Tab. 5 shows the ablation study of our proposed method compared with the baseline (QDrop) [44]. We quantize ResNet-18 and MobileNetV2 as examples to analyze the effects of PD and DC on the overall method. We evaluate our method on W2A2 and W4A2, which can reveal the impact of each component. PD-only optimizes the quantization parameters (activation scaling factors and rounding values) by only PD loss. However, there is a huge drop in results because of over-fitting, as analyzed in Sec. 3.2. We introduce Regularization to compensate for the performance loss of PD-only and achieve higher accuracy than the baseline.

Since QDrop optimizes the quantization parameters by computing the difference between the activation of the quantized block and the FP block. So our proposed DC can be well applied to this baseline method to alleviate the overfitting problem. Seeing the effect of QDrop+DC, it can improve by 0.9% at W2A2 compared with the baseline. In our PD-Quant, the introduction of DC can also further improve the performance by around 0.7% in MobileNetV2 at the W2A2 setting.

The impact of introducing prediction difference can be seen from PD+Reg. PD loss with regularization can improve the accuracy of ResNet-18 by 1.68% and MobileNetV2 by 3.21% compared with the baseline at the W2A2 bit setting. In W4A2 quantization, the performance of our method is also better than the baseline.

| Model | ResNet-18 | | MobileNetV2 | |
|-----------------|--------------|--------------|--------------|--------------|
| | W2A2 | W4A2 | W2A2 | W4A2 |
| QDrop | 51.42 | 57.56 | 10.28 | 17.30 |
| QDrop+DC | 52.32 | 57.77 | 10.38 | 17.58 |
| PD-only | 1.07 | 51.32 | 7.01 | 13.59 |
| PD+Reg | 52.74 | 58.17 | 13.49 | 20.05 |
| PD-Quant | 53.08 | 58.65 | 14.17 | 20.40 |

Table 5. Ablation study (top-1 accuracy(%)) on validation set for our proposed method. QDrop is the baseline method. QDrop+DC means to introduce DC in QDrop. PD-only means optimizing quantization parameters by only PD. Reg means regularization. PD-Quant is our proposed method, including PD, Reg, and DC for optimizing both activation scaling factors and rounding values.

4.4. Hyperparameters Analysis

As described in Eq. (5), there is a hyper-parameter λ_r in the optimization objective. λ_r limits the strength of the regularization constraint. We analyze how it affects quantization performance on ResNet-18 and MobileNetV2 with W2A2 quantization. According to Tab. 6, we found that different networks have different optimal λ_r , but our method is not very sensitive to it. Results show that the proposed method can achieve a steady accuracy without tedious hyperparameter tuning.

| Model | λ_r | | | | |
|-------------|-------------|-------|-------|-------|-------|
| | 0.05 | 0.1 | 0.2 | 0.5 | 1 |
| ResNet-18 | 52.43 | 52.60 | 52.74 | 52.58 | 52.55 |
| MobileNetV2 | 13.48 | 13.49 | 13.03 | 13.37 | 12.38 |

Table 6. Hyperparameter analysis for λ_r at W2A2 setting.

| Model | λ_c | | | | |
|-------------|-------------|-------|-------|-------|-------|
| | 0 | 0.001 | 0.005 | 0.01 | 0.02 |
| ResNet-18 | 52.74 | 53.04 | 52.87 | 52.89 | 53.08 |
| MobileNetV2 | 13.49 | 13.77 | 14.17 | 13.24 | 13.09 |

Table 7. Hyperparameter analysis for λ_c at W2A2 setting.

Moreover, we also conduct experiments to analyze the hyperparameter in Distribution Correction. As illustrated in Eq. (6), λ_c controls the degree of correction. We demonstrate how λ_c affects model performance in Tab. 7, where $\lambda_c = 0$ means no DC.

Although there is some effect when λ_c changes, our DC method can also get a stable level of accuracy without hand-crafted hyper-parameter adjusting.

5. Discussion

In this section, we will first discuss why we do not determine weight quantization scaling factors by PD loss. Then the reasons for choosing KL to calculate the prediction difference will be explained. Finally, we will analyze the limitations of PD-Quant.

5.1. Weight Quantization Scaling Factor With PD Loss

We didn't use PD loss to find weight quantization scaling factors (S_w) in Sec. 3.1 because weight quantization is different from activation quantization. Current methods mostly quantize weights per channel, resulting in a much larger search space than activation quantization. Besides, our experiments in Tab. 8 show that the weight scaling factors are not sensitive to the introduction of PD information.

| Metric | W2A32 | | W32A2 | |
|------------------|--------------|------------|--------------|--------------|
| | First layer | All layers | First layer | All layers |
| MSE ^l | 11.09 | 0.09 | 56.10 | 23.38 |
| PD ^g | 10.98 | 0.10 | 56.83 | 28.66 |

Table 8. We quantize only weights/activations for ResNet-18 to investigate how PD loss affects weight/activation quantization scaling factors. The results listed include only quantizing the first layer of the model and quantizing all layers.

We evaluate how PD loss affects activation scaling factors and weight scaling factors in Tab. 8. To avoid the influence of other quantization parameters, we optimize only scaling factors in the experiments. As can be seen, all existing methods fall crash, including PD loss, when we quantize the weights with low bits. To better understand the impact of our method on weight quantization, we conducted an experiment on the first layer only, referred to as "First Layer" for the rest of the paper. Unlike the previous experiments, the First Layer did not set the first and last layer to 8 bit, and we set per-tensor weight quantization in the First Layer of W2A32. Our method can effectively search for the optimal scaling factor during activation quantization at low bits, but not for weight quantization. We think this is because the parameter space of weights is deterministic, while the parameter space of activations varies with different input image.

5.2. Determination of PD loss

In Tab. 9, we have compared different global metrics for searching the activation quantization scaling factors. MSE^g and Cosine^g measure the difference between the quantized prediction and FP prediction by MSE distance and cosine

distance. KL^g consider the KL divergence [17] of the two models' prediction. As can be seen in Tab. 9, KL performs best among all metrics. We believe this is because our quantization process can be seen as a knowledge distillation [11, 17], where the quantized model is the student and the FP model is the teacher.

| Model | ResNet-18 | | RegNet-600M | |
|-----------------------|--------------|--------------|-------------|-------------|
| | W8A2 | W4A2 | W8A2 | W4A2 |
| MSE ^g | 27.06 | 11.13 | 3.70 | 1.64 |
| Cosine ^g | 27.28 | 7.12 | 5.93 | 3.06 |
| KL^g | 28.41 | 12.27 | 7.47 | 3.17 |

Table 9. Metric test (top-1 accuracy(%)) on validation set for activation scaling factors. ^g means calculating the difference between predictions.

5.3. Limitations

PD-Quant requires additional computation and fine-tuning compared to the baseline method QDROP, resulting in increased time cost. The time cost of quantization is shown in Tab. 10. While training time consumption is a limitation of PD-Quant, our method's additional time cost is acceptable when compared to quantization-aware training (QAT) methods. For example, LSQ [8], a QAT method, takes 120 hours to train for 90 epochs at ResNet-18, while PD-Quant only needs 1 hour.

| Method | ResNet-18 | MobileNetV2 | RegNetX-600MF |
|--------|-----------|-------------|---------------|
| QDrop | 0.43h | 0.93h | 0.89h |
| PD | 0.91h | 2.26h | 2.37h |
| PD+DC | 1.11h | 2.68h | 2.75h |

Table 10. Time cost comparison. (one Nvidia RTX A6000)

6. Conclusion

In this work, we first observed that it was very beneficial to introduce PD information when optimizing the activation scaling factors at low-bit. Based on this observation, we proposed PD-Quant, an effective method for post-training quantization. When optimizing activation scaling factors and rounding values, we discovered that our proposed method could also improve current methods. In addition, we found that over-fitting is a factor that leads to performance degradation. To respond to this problem, we propose a method to correct the calibration set distribution to improve model generalizability. The experimental results show that PD-Quant is very effective at low bits.

Acknowledgements

This work was supported by National Natural Science Foundation of China (NSFC No. 61733007).

References

- [1] Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional networks for rapid-deployment. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 2, 6
- [2] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 2
- [3] Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Zeroq: A novel zero shot quantization framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13169–13178, 2020. 2
- [4] Zhaowei Cai, Xiaodong He, Jian Sun, and Nuno Vasconcelos. Deep learning with low precision by half-wave gaussian quantization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5918–5926, 2017. 2
- [5] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*, 2018. 2
- [6] Kanghyun Choi, Hye Yoon Lee, Deokki Hong, Joonsang Yu, Noseong Park, Youngsok Kim, and Jinho Lee. It's all in the teacher: Zero-shot quantization brought closer to the teacher. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8311–8321, 2022. 2
- [7] Yoni Choukroun, Eli Kravchik, Fan Yang, and Pavel Kisilev. Low-bit quantization of neural networks for efficient inference. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3009–3018. IEEE, 2019. 1, 2, 3
- [8] Steven K Esser, Jeffrey L McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. *arXiv preprint arXiv:1902.08153*, 2019. 2, 8
- [9] Jun Fang, Ali Shafiee, Hamzah Abdel-Aziz, David Thorsley, Georgios Georgiadis, and Joseph H Hassoun. Post-training piecewise linear quantization for deep neural networks. In *European Conference on Computer Vision*, pages 69–86. Springer, 2020. 2
- [10] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. *arXiv preprint arXiv:2103.13630*, 2021. 1
- [11] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021. 2, 8
- [12] S Han, H Mao, and WJ Dally. Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015. 1
- [13] Tiantian Han, Dong Li, Ji Liu, Lu Tian, and Yi Shan. Improving low-precision network quantization via bin regularization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5261–5270, 2021. 2
- [14] Matan Haroush, Itay Hubara, Elad Hoffer, and Daniel Soudry. The knowledge within: Methods for data-free model compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8494–8502, 2020. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4, 6
- [16] Alex Hernández-García and Peter König. Data augmentation instead of explicit regularization. *arXiv preprint arXiv:1806.03852*, 2018. 4
- [17] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 2, 4, 8
- [18] Itay Hubara, Yury Nahshan, Yair Hanani, Ron Banner, and Daniel Soudry. Accurate post training quantization with small calibration sets. In *International Conference on Machine Learning*, pages 4466–4475. PMLR, 2021. 1
- [19] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2704–2713, 2018. 2
- [20] Yongkweon Jeon, Chungman Lee, Eulrang Cho, and Yeonju Ro. Mr. biq: Post-training non-uniform quantization based on minimizing the reconstruction error. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12329–12338, 2022. 2, 4, 7
- [21] Sangil Jung, Changyong Son, Seohyung Lee, Jinwoo Son, Jae-Joon Han, Youngjun Kwak, Sung Ju Hwang, and Changkyu Choi. Learning to quantize deep networks by optimizing quantization intervals with task loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4350–4359, 2019. 2
- [22] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018. 1
- [23] Jun Haeng Lee, Sangwon Ha, Saerom Choi, Won-Jo Lee, and Seungwon Lee. Quantization for rapid deployment of deep neural networks. *arXiv preprint arXiv:1810.05488*, 2018. 1
- [24] Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. Brecq: Pushing the limit of post-training quantization by block reconstruction. *arXiv preprint arXiv:2102.05426*, 2021. 1, 2, 4, 6, 7
- [25] Zhengyi Li, Cong Guo, Zhanda Zhu, Yangjie Zhou, Yuxian Qiu, Xiaotian Gao, Jingwen Leng, and Minyi Guo. Efficient activation quantization via adaptive rounding border for post-training quantization. *arXiv preprint arXiv:2208.11945*, 2022. 1

- [26] Zhengyi Li, Cong Guo, Zhanda Zhu, Yangjie Zhou, Yuxian Qiu, Xiaotian Gao, Jingwen Leng, and Minyi Guo. Efficient activation quantization via adaptive rounding border for post-training quantization. *arXiv preprint arXiv:2208.11945*, 2022. 2
- [27] Tailin Liang, John Glossner, Lei Wang, Shaobo Shi, and Xiaotong Zhang. Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing*, 461:370–403, 2021. 1
- [28] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270*, 2018. 2
- [29] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, pages 7197–7206. PMLR, 2020. 2, 4, 6
- [30] Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021. 1
- [31] Yury Nahshan, Brian Chmiel, Chaim Baskin, Evgenii Zheltonozhskii, Ron Banner, Alex M Bronstein, and Avi Mendelson. Loss aware post-training quantization. *Machine Learning*, 110(11):3245–3262, 2021. 6
- [32] Eunhyeok Park, Sungjoo Yoo, and Peter Vajda. Value-aware quantization for training and inference of neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 580–595, 2018. 1
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6
- [34] Tomaso Poggio, Kenji Kawaguchi, Qianli Liao, Brando Miranda, Lorenzo Rosasco, Xavier Boix, Jack Hidary, and Hrushikesh Mhaskar. Theory of deep learning iii: explaining the non-overfitting puzzle. *arXiv preprint arXiv:1801.00173*, 2017. 4
- [35] Antonio Polino, Razvan Pascanu, and Dan Alistarh. Model compression via distillation and quantization. *arXiv preprint arXiv:1802.05668*, 2018. 2
- [36] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10428–10436, 2020. 4, 6
- [37] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 4, 6
- [38] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 6
- [39] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2820–2828, 2019. 6
- [40] Changbao Wang, DanDan Zheng, Yuanliu Liu, and Liang Li. Leveraging inter-layer dependency for post-training quantization. In *Advances in Neural Information Processing Systems*. 2
- [41] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8612–8620, 2019. 2
- [42] Lin Wang and Kuk-Jin Yoon. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [43] Peisong Wang, Qiang Chen, Xiangyu He, and Jian Cheng. Towards accurate post-training network quantization via bit-split and stitching. In *International Conference on Machine Learning*, pages 9847–9856. PMLR, 2020. 2, 6
- [44] Xiuying Wei, Ruihao Gong, Yuhang Li, Xianglong Liu, and Fengwei Yu. Qdrop: Randomly dropping quantization for extremely low-bit post-training quantization. *arXiv preprint arXiv:2203.05740*, 2022. 1, 2, 4, 5, 6, 7
- [45] Di Wu, Qi Tang, Yongle Zhao, Ming Zhang, Ying Fu, and Debing Zhang. Easyquant: Post-training quantization via scale optimization. *arXiv preprint arXiv:2006.16669*, 2020. 1, 2, 3
- [46] Shoukai Xu, Haokun Li, Bohan Zhuang, Jing Liu, Jiezhong Cao, Chuangrun Liang, and Mingkui Tan. Generative low-bitwidth data free quantization. In *European Conference on Computer Vision*, pages 1–17. Springer, 2020. 2
- [47] Jiwei Yang, Xu Shen, Jun Xing, Xinmei Tian, Houqiang Li, Bing Deng, Jianqiang Huang, and Xian-sheng Hua. Quantization networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7308–7316, 2019. 2
- [48] Hongyi Yao, Pu Li, Jian Cao, Xiangcheng Liu, Chenying Xie, and Bingzhang Wang. Rapq: Rescuing accuracy for power-of-two low-bit post-training quantization. *arXiv preprint arXiv:2204.12322*, 2022. 4, 7
- [49] Dongqing Zhang, Jiaolong Yang, Dongqiangzi Ye, and Gang Hua. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 365–382, 2018. 2
- [50] Xiangguo Zhang, Haotong Qin, Yifu Ding, Ruihao Gong, Qinghua Yan, Renshuai Tao, Yuhang Li, Fengwei Yu, and Xianglong Liu. Diversifying sample generation for accurate data-free quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15658–15667, 2021. 2
- [51] Ritchie Zhao, Yuwei Hu, Jordan Dotzel, Chris De Sa, and Zhiru Zhang. Improving neural network quantization without retraining using outlier channel splitting. In *International*

conference on machine learning, pages 7543–7552. PMLR, 2019. [2](#)

- [52] Yunshan Zhong, Mingbao Lin, Gongrui Nan, Jianzhuang Liu, Baochang Zhang, Yonghong Tian, and Rongrong Ji. Intraq: Learning synthetic images with intra-class heterogeneity for zero-shot network quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12339–12348, 2022. [2](#)
- [53] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016. [2](#)
- [54] Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression. *arXiv preprint arXiv:1710.01878*, 2017. [2](#)